

523,517

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau




(43) International Publication Date
18 March 2004 (18.03.2004)

PCT

(10) International Publication Number
WO 2004/023312 A1

- (51) International Patent Classification⁷: **G06F 12/08**, 12/06
- (74) Agents: **WATKIN, Timothy, Lawrence, Harvey et al.**; Lloyd Wise, Tanjong Pagar, P.O. Box 636, Singapore 910816 (SG).
- (21) International Application Number: **PCT/SG2002/000176**
- (81) Designated States (*national*): CN, IN, US, ZW.
- (22) International Filing Date: 5 August 2002 (05.08.2002)
- (84) Designated States (*regional*): European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR).
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant (*for all designated States except US*): **INFINEON TECHNOLOGIES AG** [DE/DE]; St.-Martin-Strasse 53, 81669 München (DE).
- Declaration under Rule 4.17:
— of inventorship (Rule 4.17(iv)) for US only
- (72) Inventors; and
- Published:
— with international search report
- (75) Inventors/Applicants (*for US only*): **KAMIKO, Taro** [JP/SG]; Block 130 Bishan Street 12 #06-249, Singapore 570130 (SG). **PANDEY, Pramod** [IN/SG]; Block 226 Pasir Ris Street 21 #11-72, Singapore 510226 (SG).
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHODS FOR OPERATING A CPU HAVING AN INTERNAL DATA CACHE



(57) Abstract: A CPU 3 having a processor 1 and an internal data cache 7 IS operated in combination with a dummy interface 13 which simulates the existence of an external memory 17 having the same address space as the cache memory 7 but which does not store data written to it. In this way, a conventional CPU can be operated without read/write access to an external memory in respect of at least part of its memory address space, and therefore with a higher performance resulting from faster memory access and reduced external memory requirements. The CPU 3 may be one of a set of CPU chips 20, 21 in a data processing system, one or more of those chips 20 optionally having read/write access to an external memory 23.

WO 2004/023312 A1

Methods for operating a CPU having an internal data cacheField of the invention

The present invention relates to methods for using a CPU cache memory, to chipsets which perform the method, and to devices incorporating those
5 chipsets.

Background of the Invention

Conventionally, a CPU integrated circuit which is a CPU has access to two forms of memory: a first memory ("data cache") within the CPU, and an
10 external memory stored on one or more memory chips. The processor of the CPU is able to access the data cache more quickly than the external memory, but the capacity of the data cache is smaller than that of the external memory, and for many calculations the CPU requires the full capacity of the external memory. For this reason, the data cache is arranged to store a duplicate of
15 the data in the external memory. The processor can then access the data cache when it requires the duplicated data, and the external memory when it requires data which is not duplicated in the internal memory.

The configuration is illustrated in Fig. 1 which shows that when the processor
20 1 of the CPU 3 needs to access data with a given address it does so via a cache controller 5 which determines whether the data with that address is stored in the data cache 7. If so, the data is read from there. If not, the data is read from external memory 9 (which may take the form of one or more memory chips) using a system bus 11. Once the CPU 3 reads data from the
25 external memory 9, the data will be mirrored in the data cache 7, and the next read access from the same location in the external memory 9 will be skipped and instead the data will be taken from the data cache 7. When the processor writes data into the data cache 7 it is not normally copied into the external

memory 9 at once, but instead written into the external memory 9 when the corresponding data in the data cache 7 is going to be replaced. Note that the system of Fig. 1 normally includes an address decoder (not shown) which is involved in passing the signals between the CPU 3 and the external memory 9. This address decoder (not shown) is employed because the external memory 9 may be partitioned, and the address decoder defines a mapping between the address in the CPU memory space and an address in the external memory 9. For example, when the CPU 3 issues a write instruction the decoder converts the output address of the CPU 3 in the memory space of the CPU into the address of the corresponding location in the partitioned external memory 9.

The memory address map of the system of Fig. 1 in the memory space of the CPU is illustrated in Fig. 2. All of this memory corresponds to addresses in the external memory 9 (under the mapping defined by the decoder). The address space includes a "memory" section 4 which the processor 1 employs as the read/write memory for performing its operations (varying portions of this memory space are mirrored in the data cache 7), and a "reserved" section 6 which is an area of memory which is not in fact required for performing calculations.

Summary of the invention

The present invention aims to provide a new and useful way of accessing a memory, and a chip set which implements the method.

The present invention is based on the realisation that the processes of the computer could be implemented with higher performance if only the memory stored in the data cache is used. For example, in a system in which there are several processors, performance would be improved if one or more of those

processors used only their respective internal data caches. This higher performance may result from either or both of faster memory access and reduced external memory requirements.

5 In general terms, the present invention proposes that a CPU having an internal data cache is operated in combination with a dummy interface which simulates the existence of an external memory having the same address space as the cache memory but which does not store data written to it. In this way, a CPU according to present designs can be operated without the ability to store data in an external memory in respect of at least part of its memory
10 address space.

Specifically, a first expression of the present invention is a data processing system having:

at least one processor chip including a processor unit and an internal data cache, and

15 a dummy interface which receives data written to it by the processor chip, and discards it.

The processor(s) can also be connected, e.g. via the dummy interface, to an external memory chip which provides data to initialise the data cache.

20 The at least one processor chip may be one or more of a plurality of processor chips which also includes at least one processor chip having access to an external memory store.

A second expression of the invention is a method of operating a processing chip having a processor, an internal data cache and a cache controller for transmitting write instructions out of the integrated circuit for updating an
25 external memory, the method including discarding the write instructions and

arranging for the program code operated by the processor to require only the data cache as memory.

Brief description of the figures

An embodiment of the invention will now be described for the sake of
5 illustration only with reference to the following figures, in which:

Fig. 1 shows schematically a known data processing system;

Fig. 2 shows the address map for the system of Fig. 1;

Fig. 3 shows schematically a data processing system according to the present
10 invention;

Fig. 4 shows an external memory simulated by the system of Fig. 3;

Fig. 5 shows the address map for the system of Fig. 3; and

Fig. 6 shows a system incorporating an arrangement as shown in Fig. 3.

Detailed Description Of The Embodiments

15 The embodiment is illustrated in Fig. 3, in which components which may be identical to those shown in Fig. 1 are illustrated by the same reference numerals. As shown in Fig. 3, the data processing system includes a CPU processor chip 3 having an internal processor 1, a cache controller 5 and an internal data cache 7 connected via the system bus 11 to a dummy interface
20 13. The dummy interface 13 is coupled to a memory 19, and to an address decoder 15.

In an initialisation stage, the dummy interface 13 reads data (including program code) which is required to initialise the CPU 3 from the memory 19 into part of the data cache 7.

Subsequently, in operation, the processor 1 operates exactly as in present systems. However, the program code which it runs (i.e. the program code uploaded from the memory 19 during the initialisation) is such that the amount of memory required by the processor 1 is provided by the data cache 7. The cache controller 5 operates in the conventional manner, performing write operations to the data cache 7 and attempting to keep an (in fact non-existent) external memory 17 updated to mirror what is in the data cache 7 by transmitting write operations to the system bus 11. These write instructions are passed to the dummy interface 13 which simply discards the data it receives.

The address decoder 15 operates in a way similar to that of the prior art decoder discussed above, but with the difference that it never directs read/write commands from the CPU to the external memory 9. Write commands from the CPU 3 are simply discarded by the dummy interface 13 and the decoder 15. In the case that the dummy interface receives a read command from the CPU 3 (as in some embodiments it may do during the initiation procedure), the dummy interface transmits an output which is not dependent on the data which has previously been transmitted to it, e.g. it may always output "0". This feature may be used during an initialisation of the data cache 7 in which portions of the memory which are not copied from the external memory 19 may be initialised to zero using values read from the dummy interface 13.

25

We now analyse the situation from the perspective of the CPU 3. The operation of the CPU is identical to what it would be if the dummy interface were not just discarding data, but instead was a normal interface performing read/write operations to a memory 17, operating as shown in Fig. 3: that is, not storing ("N/C") any data written to it, and always outputting the value "0" in

30

any read operation. In fact, this memory 17 does not exist, but the CPU 3 has no way of "knowing" this.

Since the program code is such that the only memory required is provided by the data cache 7, the processor 1 operates in exactly the same manner as it would in the system of Fig. 1 running the same code. However, the system of Fig. 3 has the advantage that less memory is required, since the external memory 19 of Fig. 3 is only required to store the initialisation data and may be smaller than the memory 9 of Fig. 1.

Turning to Fig. 5, the address map for the system is shown. In this case, the memory space of the CPU 3 is entirely defined by the data cache 7. A portion 14 of the data cache stores data which was read from the external memory 19 during initialisation, and this section is not overwritten during the operation of the processor 1. A second portion 16 of the data cache 7 is the read/write area which the processor 1 uses to perform its calculations. A portion 18 of data cache 7 is "reserved", i.e. unused in the operations of the processor 1.

Note that the CPU 3 may be one of a plurality of processing units which are part of the data processing system, such as the shown in Fig. 6 having a master processing unit 20 and a plurality of slave processing units 21. In this case, one or more of the CPU processing units (e.g. the master processing unit 20 only) may be provided with an external memory 23 which is kept updated to mirror that processing unit's internal data cache by a conventional method such as that described above in relation to Fig. 1. Other of the CPU processing units (e.g. some or all of the slave processing units 21) may be as discussed above in relation to Fig. 3, i.e. not being able to store data an external memory but instead performing processing only using their internal data caches and transmitting their write instructions to dummy interfaces 25 which discard them. The dummy interfaces initialise the slave processing unit 21 using one or more memory chips 29 which do not store data during the

subsequent processing operation. In this way, the overall performance of the system is improved, since fewer read/write memories 23 are required. At the same time, the processing speed of the system is improved since many of the processors (e.g. the slave processing units 21) never need to wait to obtain data from external memories.

In such a system, the only processor chip(s) 20 which are provided with the capacity to read/write data to/from an external memory 23 are the one(s) which require a read/write memory larger than their internal data cache to perform their tasks, and the program code which the full set of processing chips 20, 21 operate is written such that all tasks which require a memory larger than the internal data cache of the CPUs are allotted to those CPU(s) 20 which have access to an external memory.

Claims

1. A data processing system having:

at least one processor chip including a processor unit and an internal data cache, and

5 a dummy interface which receives data written to it by the processor chip, and discards it.
2. A data processing system according to claim 1 in which the dummy interface is coupled to a memory, the dummy interface passing data to the processor chip during initialisation.
- 10 3. A data processing system according to claim 1 further including one or more further processing chips which have read/write access to external memory.
4. A method of operating a processing chip having a processor, an internal data cache and a cache controller for transmitting write instructions
15 out of the integrated circuit for updating an external memory, the method including discarding the write instructions and arranging for the program code operated by the processor to require only the data cache as memory.

FIG 1

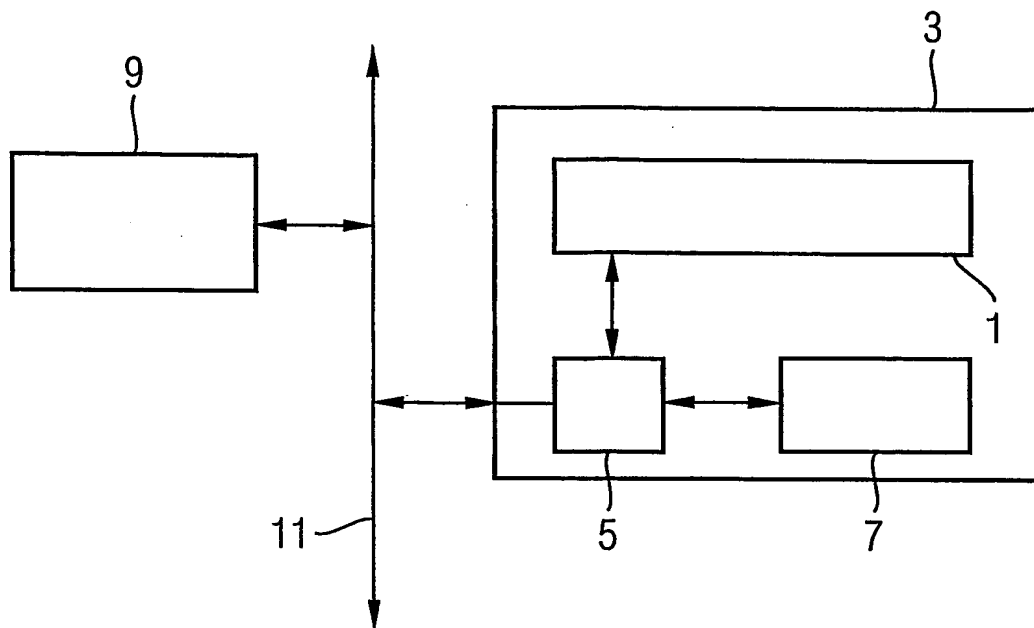


FIG 2




FIG 3




FIG 4




FIG 5





FIG 6



INTERNATIONAL SEARCH REPORT

International Application No.

PCT/JP02/00176

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F12/08 G06F12/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 80 00383 A (HIGGINS L ; SQUARE D CO (GB)) 6 March 1980 (1980-03-06)	1, 2, 4
Y	page 2, line 20 - page 4, line 14; figure 1	3
Y	US 5 692 149 A (LEE GYUNGHO) 25 November 1997 (1997-11-25) column 4, line 34 - column 5, line 3; figure 1	3

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

S document member of the same patent family

Date of the actual completion of the international search

10 April 2003

Date of mailing of the international search report

28/04/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Ledrut, P

INTERNATIONAL SEARCH REPORT

International Application No
PCT/SG 02/00176

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 8000383	A	06-03-1980	
		AU 536458 B2	10-05-1984
		AU 4979479 A	06-03-1980
		CA 1139891 A1	18-01-1983
		DE 2967707 D1	05-11-1992
		DK 153580 A	10-04-1980
		EP 0018404 A1	12-11-1980
		ES 483324 A1	01-09-1980
		WO 8000383 A1	06-03-1980
		IE 48797 B1	15-05-1985
		IT 1119151 B	03-03-1986
		JP 55500695 T	25-09-1980
		ZA 7904177 A	29-04-1981
US 5692149	A	25-11-1997	NONE